Consciousness and Cognition 126 (2024) 103783

Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/yccog

Full Length Article

Valence in perception: Are affective valence and visual brightness integral dimensions in visual experience?

Hilla Jacobson [a,*], Zohar Ongil [b], Daniel Algom [b,c], Marius Usher [b]

[a] The Hebrew University of Jerusalem, Israel
[b] Tel-Aviv University, Israel
[c] Achva Academic College, Israel

ARTICLE INFO

Keywords:
Valence
Affect
Affect-perception interactions
Separable vs. integral dimensions
Garner interference
Consciousness
Phenomenal character

ABSTRACT

A fundamental question in the domain of affect and conscious perception is whether the former can impact the latter. Traditionally, perception and affect were conceived as largely independent. Against this backdrop, it was recently argued that the affective valence of a stimulus can modulate the perceptual experience of its sensory features. An alternative hypothesis is that perceptual experiences have a valenced aspect over and above their sensory aspects, with these two aspects interacting and comprising integral perceptual dimensions. To test this, we carried out two experiments deploying Wendell Garner's speeded classification paradigm to decide whether visual brightness and affective valence are *separable* or *integral* dimensions. We found Garner interference, documenting that brightness and valence are integral dimensions. We did not observe effects of congruity – responses to bright positive stimuli were not faster than to bright negative stimuli — providing no support for affect induced changes in the perception of brightness.

1. Introduction

What is the relation between (conscious) perception and affect? Does affect impact how we consciously perceive the world, and if it does, in what way? Can things (literally) *look* good or bad? Recently, there has been growing interest in the extent to which the cognitive and neural mechanisms underlying perception and affect are distinct or integrated (Phelps et al., 2006; Meier et al., 2007; Pessoa, 2008; Barrett and Bar, 2009; Stefanucci & Storbeck, 2009; Banerjee et al., 2012; Song et al., 2012; Pourtois et al., 2013; Ferneyhough et al., 2013; Topolinksi et al., 2015; Vuilleumier, 2015; Carrasco & Barbot, 2018; 2019; Lindell et al., 2022). Whereas in the past, these two basic mental capacities were thought to be distinct and were studied separately, this is no longer the case. During the past decades, converging evidence has accrued that demonstrates the various ways in which perceptual and affective mechanisms are strongly interlinked (Pessoa, 2008; Duncan & Barett, 2007). Thus, Pessoa (2008) argues that emotion and perception (as well as cognition more generally) are only minimally decomposable, and Duncan and Barrett (2007) assert that the affect-perception distinction “is not respected in the human brain” (p. 1).

Different lines of research support the view that affective processes occur automatically and often precede perceptual analysis (Bargh et al., 1992). Moreover, the two processes likely interact bidirectionally. Let us begin with the link from perception to affect. First, many studies have shown fast affective evaluations of visual stimuli, often preceding perceptual analysis (Bargh et al., 1992

* Corresponding author at: Department of Philosophy and Department of Cognitive and Brain Sciences, Israel.
E-mail address: hilla.jacobson@mail.huji.ac.il (H. Jacobson).

https://doi.org/10.1016/j.concog.2024.103783
Received 21 April 2024; Received in revised form 31 October 2024; Accepted 31 October 2024

Handy et al., 2010; Kuhbandner et al., 2011). Second, it has been shown that even 'neutral' visual stimuli can trigger rapid online affective responses as a result of fluency manipulations that affect the ease of perceptual processing (Reber et al., 1998; Reber et al., 2004; Topolinski et al., 2015; Lindell et al. 2022). For example, Topolinski et al. (2015) have shown positive affect for stimuli that allowed (rather than obstructed) Gestalt completion, indicating that affect is a consequence of early perceptual processes. In that study, several measures were taken to support the hypothesis that the affective responses were triggered in early processing stages and are thus *online-components of perception*. The duration of stimulus presentation was limited to 100 ms (and even 25 ms in some experiments) and the affective responses were probed not only by verbal reports of stimulus-liking but also by measuring immediate physiological responses (by facial EMG). The very brief presentations, which still allowed perception of the stimuli but prevented conscious insight into the underlying manipulations, ruled out more conscious, strategic, or semantic processes as mediators of the effects. The results of this study are consistent with an account of emotion in perception that proposes that visual stimuli undergo various evaluation checks at different levels of processing (e.g., Leventhal and Scherer, 1987) and indicates that affect is modulated by the success of visual process (e.g., visual disambiguation).

Other studies have focused on interactions in the opposite direction – i.e., from affect to perception – illustrating the influence of affective, valenced features of the stimulus on perceptual processing. One of the best known such demonstrations is the Emotional Stroop Effect (ESE; Algom et al., 2004; McKenna, 1986; Williams et al., 1996). This effect arises when participants name the color of printed words while ignoring the meaning of the words themselves. The words come from two categories, emotionally neutral or negative. The ESE documents the result that it takes longer to name the color of negative words than to name that of neutral words. The ESE is likely mediated by a difference in success of selective attention to the task relevant colors across emotions (Mama et al., 2013). It is possible, though, that the ESE takes place during a *late stage* of the perceptual process (e.g., during *response-selection*, see Cohen et al., 1990).

Other lines of research, however, have shown that rapidly encoded affect (Handy et al., 2010) can enhance processing efficiency and modulate *early stages* of perceptual processes, such as object recognition. These studies also illustrate that the two directions of influence often take place together, in an interacting loop. For example, the operation of affective mechanisms improves perceptual performance, resulting in enhanced ability to attend, search, and track affectively salient visual stimuli (e.g., Bargh,1992; Bargh et al., 1992; Chen & Bargh, 1999; Phelps et al., 2006, Domínguez-Borràs et al., 2017). Moreover, Barrett and Bar have contended that we see objects "with feeling" (Barrett and Bar, 2009, p. 1325). Their argument appeals to neural circuitry and to the notion of affective predictions: the brain routinely and automatically makes affective predictions during visual processing, based on representations of the affective impact of similar stimuli in the past. These predictions, first formed only milliseconds after visual stimuli register on the retina, result in affective responses signaling objects' value, and these responses are *integral to seeing* the object as what it is.[1] According to Barrett and Bar's hypothesis, "people do not wait to evaluate an object for its personal significance until after they know what the object is. Rather, an affective reaction is one component of the prediction that helps a person see the object in the first place" (ibid., p. 1331). It should also be mentioned that support for behavioral research that aims to establish the *early* involvement of valence in perceptual processes is also provided by the accumulation of data regarding brain mechanisms that can subserve these processes. Of special importance is evidence regarding the role of the amygdala as a hub that encodes and monitors the affective value of stimuli, and consequently sends (direct as well as indirect) feedback to sensory pathways, including the striate and extrastriate visual cortex (e.g., Ledoux, 2000; Pourtois et al., 2013; Vuilleumier, 2015).

Finally, the role of affect in perception (including vision) has been proposed to extend to daily objects (and thus beyond the 'guns and roses' type stimuli that are typically used in lab experiments). This generalization is based on the idea that daily visual objects possess 'micro-valences' (Lebrecth and Tarr, 2010; Lebrecht, 2012; Lebrecht et al., 2012; Barrett & Bar, 2009). Thus, it has been argued that valence is ubiquitous: "The majority of physical objects around us possess a subtle valence – a 'micro-valence' – which ranges in magnitude but is always present" (Lebrecht et al., 2012, p. 1). Micro-valence has been experimentally traced using both implicit and explicit measures. Thus, in the Birthday Task experiment (Lebrecht, 2012; Lebrecht & Tarr, 2010), participants were asked to select (briefly presented) objects that they would most like to keep or return as birthday gifts. In the Ranking Task (Lebrecht, 2012) participants were explicitly asked to rank-order objects along the dimension of valence (from most negative to most positive). In both experiments, participants were consistent in their selection across different trials, and the majority of participants showed strong correlations between the two measures. It has been argued, further, that the role of micro-valences in perceptual processes is similar in various respects to that of other visual features of the stimuli. Thus, Lebrecht and colleagues argue that "valence… is not a label applied after the fact to perceptual entities, but rather is an intrinsic element of visual perception with the same mental status as other object properties" (Lebrecht et al., 2012, p. 4).

2. Does affect impact on visual experience?

We can now move closer to the main concern of this study. The claim that *affect is integral to perception* (or to 'seeing') can be

[1] Barrett and Bar report that valence is rapidly encoded (as early as 80 to 130 ms after stimulus onset) in a specific region of the prefrontal cortex, the orbitofrontal cortex (OFC), that plays a crucial role in forming predictions that support object perception. Valence is consequently encoded in lower areas in the lateral occipital cortex (LOC) that are part of the "core" network for visual processing, and to which the former areas project. In particular, the encoding of valence in the OFC provides crucial evidence, as it is taken to be part of the first of two stages hypothesized to be involved in object recognition – a fast gist-based stage, which is then followed by a second slower stage of refinement. The second stage is assumed to be facilitated by the encoding of valence in the first stage.

understood in several ways. First, most of the studies just surveyed support this claim in the sense that affect is an integral part of the *perceptual process* of seeing. Perception and affect are *overlapping processes*. Affective components are triggered in immediate response to success or failure of early visual perceptual processes, and early processed affective responses influence perceptual processes. Yet, our main concern in what follows is with whether affect can be integral to perception in a related but somewhat different sense: we will examine the claim that perception and affect *share a mental state*. Specifically, we focus on the question whether affect can be an integral part of the *visual experience* of seeing, in having an impact on *the way the stimulus looks* (i.e., on the phenomenal character of the experience). In other words, does affect influence the character of our visual experiences? In this context we will consider two further (compatible) hypotheseses: first, affect can make a stimulus visually appear brighter/darker than another stimulus; second, affect can actually make a stimulus visually appear good/bad (see section 2.2 for a discussion of what this may involve). Either one of these hypotheses would imply that affect is integral to the mental state of seeing in having an impact on the way the stimulus looks, but only according to the second hypothesis does the visual experience itself have an additional dimension of valence.

It should be noted that the claim that affect is integral to perceptual *processes* does not imply that it is integral to *visual experiences*: it is possible that, e.g., the successful perception of a stimulus, or its being appraised as having positive value, *leads* to a positive affective state that *co-occurs* with, but is distinct from, the *visual* experience and does not influence how this stimulus looks. Nonetheless, we believe that the former claim provides some motivation for examining the latter claim, as it shows that perceptual and affective processes interact even at early stages and that *some* affective responses to visual stimuli are not the result of *post*-perceptual appraisals. In the rest of this section we elaborate on the question of whether, and if so how, affect influences visual experiences, and present three hypotheses regarding this question.

2.1. Three hypotheses regarding the impact of affect on visual experience

We set out to study three specific hypotheses, by examining a particular case of perceiving affectively charged stimuli. According to the first hypothesis (H1), affect does not influence visual experiences. It has no impact on perceptual appearances. The seemingly 'affective glosses' on perception (e.g., the negative affect associated with seeing an angry or sad face) are the result of *post-perceptual* appraisals. Thus, *affect does not influence the way things look* (e.g., Block, 2023; Firestone & Scholl, 2016).

According to the second hypothesis (H2), things *look* different when they are appraised positively or negatively. Specifically, affect alters perceptual appearances by modulating the *sensory aspects*[2] of a visual experience. Due to connections between brain areas that encode affect and sensory information, 'core affect' can change the specific sensory features the stimuli appear to have. For example, Carrasco and Barbot (2018, 2019) argue that the appearance of a fearful face in a particular area in the visual field enhances the *perceived contrast* of subsequent Gabor patches (presented at the same location). Similarly, Meier et al. (2007) argue that positive words are seen as *brighter* than negative words, and Song et al. (2012) argue that smiling faces appear *brighter* relative to neutral faces (but see the criticisms in Firestone & Scholl, 2016). Thus, according to H2, the affective value of the stimuli makes a difference to the way they look, which consists in a difference in *sensory* appearance.

The third hypothesis (H3) agrees with H2 in holding that affect has effects on visual experiences: affect impacts conscious perception and makes a difference to the way things look. Yet, these two hypotheses differ crucially. As we have seen, H2 proposes that the influence of affect on the phenomenal character of conscious perception consist in modulation of *sensory aspects*. H3, though, suggests that perceptual experiences have a *valenced aspect* over and above their sensory aspects (Jacobson, 2021 and forthcoming). That is, whereas according to H2, the angry face looks *darker* relative to a neutral face, according to H3, the face actually looks *bad* or *repellent*. (Note that H2 and H3 are not mutually exclusive: an angry face may look both darker and bad.) Thus, according to H3, the visual experience itself has an affective dimension – it is 'affectively-loaded' or 'intrinsically valenced.'[3] The valence is a constitutive, integral, aspect of the visual experience.[4] The sensory and valenced aspects are *bound* and *phenomenally integrated* (for elaboration, see 2.2 below), and this integration expresses itself at the level of appearances, so that the visual experiences themselves have hedonic tones and are phenomenally 'positive/negative.' In enjoying (or suffering, as the case may be) such hedonically charged visual experiences, things literally *look good/bad* to us.

It is worth noting that the idea that *some* (interoceptive as well as exteroceptive) experiences intrinsically involve affective components enjoys some initial support from phenomenological observations. The obvious example is pain, which is experienced as inherently unpleasant. Consider also gustatory and olfactory experiences. For the coffee lover, the coffee both *tastes and smells good* – the deliciousness certainly appears to be an integral part of the gustatory/olfactory experience. Similarly, the perceptual experiences of smelling and tasting rotten food can be quite awful; introspectively, such experiences resemble pains (and various unpleasant tactile sensations) in that unpleasantness seems to be a fundamental aspect of what it is like to have them. *Prima facie*, then, the idea that the

[2] By 'sensory aspects' we simply mean the standard, (non-valenced) aspects of perceptual experiences, which carry information about the (non-evaluative) objective features of stimuli in a specific *sensory-modality*, like color, length, or shape (including meaningful shapes like faces) in the case of vision, and saltiness or sweetness in the case of gustation. According to the orthodox view, in contrast to what we will propose below, perceptual experiences have only sensory aspects.

[3] The idea that affect can result in sensory modulation and the idea that perception itself has an affective dimension beyond its sensory dimensions are compatible in the sense that there can be effects of both kinds – different cases may exhibit different kinds of influences of affect on perception. For ease of exposition, we present two mutually exclusive hypotheses, but see discussion.

[4] Note that according to this hypothesis, the affective or valenced aspect of perceptual experiences is a simple, rudimentary, 'positive/negative' aspect, and not a full-blown emotion.

experiential world is a 'hot,' inherently affective, world rings plausible. However, such observations have little force in the case of some other sense-modalities, least of all in vision.[5] With respect to visual experiences – which are the focus of this paper – the idea of affective visual perception crucially depends on empirical, scientific investigation.

We may now summarize the difference between the three hypotheses as follows. According to H1 – the 'Valence *and* Perception Hypothesis' – the fact that a seen face has a positive or negative expression changes nothing at the level of its perceptual experience. In particular, it neither changes the perceived brightness of the face, nor does it add a valenced dimension to its perception. It can only influence post-perceptual appraisal judgments and affective states (Firestone & Scholl, 2016). According to H2 – the 'Sensory Modulation Hypothesis' – the induced valence can alter the appearance of the face by changing its perceived *sensory* features: it makes the face look brighter (for positive valence) or darker (for negative valence). Yet, H2 is compatible with the claim that the perceptual experience is *affectively* neutral – i.e., that it does not *feel* (specifically, look) good/bad. Similarly to H2, according to H3 – the 'Valence *in* Perception Hypothesis' (hereafter, VIP) – the fact that a seen face has a positive/negative expression may alter the visual appearance of the face. But it is only according to H3 that the perceptual experience itself is valenced. The perceptual experience has a valenced aspect beyond its sensory aspects: the induced valence may result in the face *looking* good/bad, where it's looking good/bad (i.e., its valence) is a further phenomenal dimension that is orthogonal to, yet integrated with, the sensory dimensions of this experience.

So far, most of the research on whether affect influence visual experiences, has focused on the contrast between H1 and H2, usually in the context of the cognitive penetrability question – namely, whether cognitive or affective states, exert direct *top-down* influences on perception (see Firestone & Scholl, 2016, for a review, and General Discussion section).[6] To the best of our knowledge, there has been no research testing H3 directly (though Barrett & Bar, 2009, and Lebrecht et al., 2012, may bear on it). Perhaps this is because the notion of a further phenomenal dimension that is integral to a perceptual experience yet orthogonal to its sensory aspects appears difficult to probe. In section 3, however, we suggest that the notion of such a further dimension can be operationalized via the Garner task, which probes for separable vs. integral dimensions. Before turning to this operationalization (section 2.2), we present a philosophical analysis that aims to clarify this notion. Empirically oriented readers (who prefer to focus on the operational and empirical components/aspects), may skip section 2.2 below and turn directly to section 3 (the Garner paradigm).

2.2. *Further explicating VIP (H3)*

Whereas the idea of sensory modulation (H2) is familiar from other contexts, the notion that perceptual experience has a further dimension – specifically, a valenced dimension (as per H3) – is less clear and poses theoretical difficulties (Jacobson, 2024). For example, what does it even mean for a perceptual experience, in itself, to be intrinsically valenced? That is, in what sense can valence be part of the identity of particular perceptual experiences *qua perceptual* experiences? Providing such a sense is necessary if the idea of 'valence *in* perception' is to be differentiated from the idea of 'valence *and* perception'– namely, the idea that seeing an angry face is merely *accompanied* by a distinct (non-perceptual) affective state (H1). This challenge is pressing, because it is most plausible that valenced aspects are not sensory aspects: plausibly, valenced aspects are not sensory looks, tastes, etc. For one thing, affective values are not objective features trackable by specific sense modalities, and they are suggested to be a-modal (see e.g., de Vignemont, 2021). Further, the difference between valences and sensory aspects is reflected at the phenomenological level (the 'level of appearances'). Even if things can look bad to us, we do not see 'badness' in quite the same sense in which we see 'redness' or 'squareness': reflecting on the large variety of things that can look bad (e.g., blood, angry faces and cockroaches), it is highly implausible that they share any distinctive, unique, appearances. As de Vignemont (ibid. p. 10) puts it, "valence does not look like anything (or it looks like too many things)." Thus, the question arises how can a particular valence be part of, e.g., a specifically *visual* perceptual experience? Similarly, how can an angry face *look* bad? It isn't clear how unpleasantness can characterize the way the face looks without being itself a look. Finally, a constraint on any account of the relation between valence and affect is that it be compatible with the phenomenon of valence-variance (VV): experiences that share their sensory aspect (S1) can have different valenced aspects along the spectrum of positive and negative valences (e.g., V_1, V_0, or V_{-1}).[7] To use a familiar gustatory example, the same food (e.g., eggplant) can taste delicious or unpalatable (for different perceivers as well as for the same perceiver at different times). VV has been shown to be a pervasive phenomenon in all sense-modalities (Birch, 1999; De Houwer et al., 2001). Yet, given that valences are not sensory aspects, there may seem to be tension between VV and the notion of 'valence *in* perception': how can valences be a necessary and integral part of the identity of perceptual experiences with specific sensory aspects, if they can vary independently of those aspects? Given that valences are not visual sensory aspects, VV seems to invite the idea that valences are 'external' to the perceptual experience (as per H1). Nonetheless, as we shall now explain, Jacobson (2024) proposes that (non-sensory) valences can be intrinsic to visual experiences

[5] This is true even of visual experiences of apparently affectively loaded objects such as snakes or bloody weapons, all the more so of apparently affectively neutral objects such as ordinary looking tables and teapots. Some negative affect may seem to be *associated* with, e.g., the visual experience of a snake, yet, introspectively, it isn't clear that this affect is *integral* to the experience (rather than being an emotional reaction that the visual experience arouses).

[6] The traditional view is that perception is encapsulated from higher-level cognition, as supported by the presence of visual illusions that are robust to knowledge of the illusion itself. The cognitive penetrability hypothesis entails a specific breakdown of the supposed modularity of perception, which involves a *top-down* modulation from cognition or affect to perception (Firestone & Scholl, 2016). As we explain later, we believe that the cognitive penetrability question is orthogonal to the question we raise in this paper. If affect impacts perception, this effect may be either bottom-up or top-down. We believe, though, that in central cases the effect is *bottom-up* (Bargh, 1992; Chen & Bargh, 1999).

[7] We denote positive and negative valences with positive/negative numerals.

notwithstanding VV. Despite being orthogonal to visual sensory aspects, valences can be intrinsic to visual experiences due to their relation to, or the particular way that they are integrated with, such aspects. Specifically, an appropriate model for VIP should take valences and sensory aspects to be (logically) independent, orthogonal, features, yet such that the phenomenal character of a given experience is a non-compositional function of both features.

A key feature of Jacobson's (2024) account of VIP is the principle she terms "*Intrinsicality as Phenomenal Uniqueness*," which is supported by its theoretical import in resolving the abovementioned theoretical difficulty (and others). The principle is to be understood as presupposing that each sensory aspects (S_1, S_2…, S_n) can be bound with different valences (V_{-n}, V_0…, V_n) and vice versa. It reads as follows: "An experience with sensory aspect S_1 has its valenced aspect V_1 intrinsically, if V_1 cannot be instanced in isolation from S_1 *while maintaining its phenomenal identity* (V_1 is 'phenomenally non-factorable' from S_1), so that when V_1 is bound with S_1 the way it's like to experience V_1 is phenomenally unique" (p. 20). A similar principle is argued to hold in the opposite direction – namely, when S_1 is bound with V_1, the way it's like to experience S_1 is phenomenally unique.

Let us explain and illustrate this principle. The principle is supposed to capture the relations between (or 'binding' of) sensory aspects (e.g., colors or shapes) and the valenced aspect that characterizes them (badness/goodness) within the same perceptual experience. It describes those relations at the phenomenological level – the level of how things appear to the perceiving subject. According to it, the binding of V_1 with S_1 and the binding of V_1 with S_2 do not share a *phenomenally* distinct component or *distinct* 'look'; rather, valences are (phenomenal) ways of having experiences with specific sensory aspects. Let us be clear: the claim is that the stimulus does *perceptually appear* valenced – it *looks* positively/negatively to some degree (the experience is pleasant/unpleasant). But at the level of appearances, the valence and the sensory aspect cannot be pulled apart and experienced as distinct, independent, features, which captures the intuitive notion that the valence appears 'fused' with the sensory aspect. At this level, 'sensory-valence binding' is 'non-compositional binding' – a kind of binding that is stronger than 'sensory-sensory' binding.

To exemplify and further clarify these ideas, it is helpful to contrast non-compositional sensory-valence binding with 'compositional' sensory-sensory binding – e.g., the binding of colors and shapes. Consider your experiences when looking at different shades that share their color – a triangle and a circle with the exact same shade of blue (Fig. 1.A).

These experiences do share a distinct phenomenal component – *with regard to color*, the (blue) triangle and the circle *look exactly the same*. Thus, the color aspect and the shape aspect are *phenomenally* decomposable. On the other hand, sensory-valence binding is different. Focusing, for illustration, on gustation (in which valence is more pronounced than in vision), sensorily different gustatory experiences, such as that of spoiled wine and rotten apples, both taste bad, but they *do not taste bad in the same way*. When 'tasting bad' is instanced in conjunction with different sensory aspects, the two instances do not share a *phenomenally* distinct valenced component – there is no *distinct* negative *valenced* aspect that *tastes exactly the same*. Rather, the result is phenomenally unique – each substance tastes bad in its own distinctive way. The sensory and valenced aspects are thus not phenomenally decomposable. Lastly, Jacobson notes that sensory-valence binding, notwithstanding its dissimilarity to color-shape binding, is similar to, e.g., the 'binding' of hue and brightness (in color perception). At the level of appearances, the relation of hue and brightness illustrates non-compositional binding. Consider experiences of two different hues that share the exact same level of brightness (Fig. 1.B). In this case, there is a commonality – the colors both look bright and their level of brightness is surely part of the way they look. However, the colors do not look bright in exactly the same way. When the same level of brightness is bound with different hues, the result is phenomenally unique, and each specific color looks bright in its own distinctive way (looking at Fig. 1A and 1B, we may note that the similarity in color 'pops out' in 1A, unlike the similarity in brightness in 1B). In this sense, hue and brightness are not decomposable phenomenal features.

The suggestion, then, is that valence is yet another *dimension* of the way things look (i.e., of *visual* phenomenal character), a

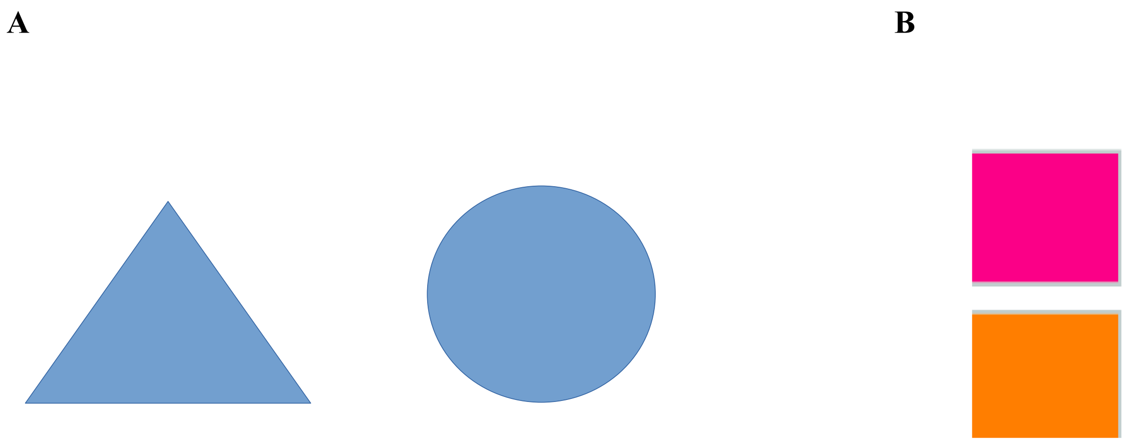

Fig. 1. An illustration of color-shape 'compositional binding' (1.A) vs. hue-brightness 'non-compositional binding' (1.B). The distinction captures two different ways in which the relations between aspects is reflected at the level of appearances. In (A), the circle and the rectangle appear exactly the same with regard to color – they share a distinct, identical appearance. In (B), although both colors look bright to the same degree, their level of brightness is not experienced (does not appear) as a distinct, identical, component; rather when the same level of brightness is 'bound' with different hues, the result is 'phenomenally unique.' The brightness changes how a hue looks, and is 'intrinsic' to a specific color look; but it has no distinct phenomenology independently of it.

dimension that is *orthogonal* to the sensory aspects with which it is bound (i.e., the same sensory aspect can be bound with different valences and vice versa). But furthermore, specific valences and specific sensory aspects cannot be instanced independently of each other while maintaining their phenomenal identity – at the level of appearances, they are strongly, non-compositionally, bound together, forming non-factorable looks. The valence changes how a sensory aspect looks, and is thereby 'intrinsic' to a specific look, yet it has no distinct phenomenology independently of the sensory aspect. Thus, the strong way in which valence is bound with sensory aspects explains how, despite being an a-modal dimension, it can be inseparable from, and integral to, *visual* experiences.

3. Operationalizing H3: The Garner paradigm

How can the Valence *in* Perception hypothesis be operationalized and empirically tested? The principle of Intrinsicality as Phenomenal Uniqueness captures and makes more precise the intuitive notion that the valence and the sensory aspects of conscious perception are phenomenally 'fused together,' that the sensory aspect has an affective gloss, or that the perception is 'affectively loaded.' It reflects the idea that, at the level of appearances, it is difficult to tease apart the valence and the sensory aspects. And the latter idea, in turn, brings us closer to home in suggesting that *it* is *impossible to attend to one aspect while ignoring the other*.

The above characterization brings to mind the notion of integral dimensions and their distinction from separable dimensions, which are empirically tested by Garner's Speeded Classification Paradigm (Garner, 1974; see also, Algom & Fitousi, 2016; Melara & Algom, 2003). As we explain below, integral dimensions (in contrast to separable dimensions) are ones that are bound in a way that precludes the possibility of drawing selective attention to one dimension of a stimulus while ignoring its other dimension. We thus propose that the Garner paradigm is a useful tool for testing H3.

Garner's paradigm provides an empirical method to decide between integral vs. separable perceptual dimensions. To illustrate, the dimensions of shape and color are found separable: One can attend to shape and ignore color and vice versa (Algom & Fitousi, 2016; Garner, 1974; Sabri, Melara, & Algom, 2001). On the other hand, color dimensions, such as hue/saturation/brightness or even height/width (for rectangles), are found to comprise integral dimensions, which means that they are encoded together in a holistic manner (Ganel & Goodale, 2014; Niv, Moran, & Algom, 2022). An illustration of how the Garner paradigm works is shown in Fig. 2. In this experimental paradigm, one presents the stimuli for categorization on one (relevant) dimension that varies from trial to trial in a random fashion. The same stimuli are presented in two types of blocks. In the baseline condition, the irrelevant dimension is kept constant across all trials (and the only variation takes place in the relevant dimension). In the filtering condition, both dimensions vary independently from trial to trial (see Fig. 2, top two rows).

The Garner interference – the difference in mean RT between filtering and baseline – is based on the idea that if two perceptual dimensions are separable, then one can selectively focus on the task-relevant dimension (A) while ignoring variations in the other (task-irrelevant) dimension (D). Selective attention will fail when the two dimensions are integral (Algom & Fitousi, 2016). In the baseline blocks, *only* the relevant dimension changes from trial to trial, making the task easy (almost not requiring selective attention). On the other hand, in the filtering blocks, both dimensions vary, and one needs to selectively attend to the relevant dimension in order to filter out irrelevant variations in the stimuli. Thus, with separable dimensions, we expect no Garner interference (as selective attention makes classification in the filtering and baseline blocks equally fast). In contrast, with integral dimensions, we expect a positive Garner interference (specifically, mean RT in Filtering > mean RT at baseline, as in Fig. 2), as the baseline conditions are easier (it is easy to ignore a constant feature of the stimulus). As mentioned, a good example of integral dimensions are width and height (for rectangles), or hue and brightness (for colors), as participants repeatedly show poorer performances when both dimensions are being changed (Filtering blocks) compared to when one of the dimensions is held constant (Baseline blocks) (Ganel & Goodale, 2014; Freud & Ganel, 2015).

In this study we applied the Garner paradigm to the dimensions of visual brightness (bright/dark) and emotional valence (positive/negative). The stimuli were smiley and frowny shapes superposed on gray circles of variable brightness. This is motivated by previous reports that smiley faces appear brighter in perception than frowny ones (Song et al., 2012; but see Firestone & Scholl, 2016). Apart from its role in assisting the current theoretical resolution, the Garner task possesses further appealing features. It is a basic and simple task entailing brief stimulus presentations and speeded responding. These features make it unlikely that strategic/semantic/ cognitive, indeed post-perceptual processes, affect the results (see Topolinski et al., 2015 on the role of short stimulus exposure; see Reber et al., 1998, on further possible effects of stimulus exposure on affect).

According to H3, valence and brightness could be considered integral dimensions, leading to a difficulty in ignoring the valence of the stimulus (its emotional aspect). It should be noted that independent of the Garner interference effect (which is a block type effect that contrasts filtering and baseline blocks), the Garner paradigm can also serve to assess the congruity effect (a type of Stroop effect; Stroop, 1935) – a within-block trial type effect that compares the RT (or accuracy) on congruent trials (bright-smiley or dark-frowny) with those on incongruent trials (bright-frowny or dark-smiley).[8] The most plausible mechanism likely to generate congruity effects in our study is the actual modulation of the sensory property (as per H2), making a gray patch appear brighter/darker when a smiley/ frowny symbol is superposed on it, as illustrated in Fig. 3.[9] The corroboration of H2 for our stimulus would require documenting

[8] The association of brightness with smiley is assumed to correspond to cultural norms that are reflected in our language and is empirically motivated by the study of Song et al. 2012.

[9] Other candidate mechanisms for the generation of congruity effects (alternative to the one we focus on here, H2) are response competition, which results from a semantic clash (Cohen et al., 1990), or a redundancy target effect caused by a race of independent channels (Eidels, 2012; for reviews see, MacLeod, 1991; Melara & Algom, 2003).

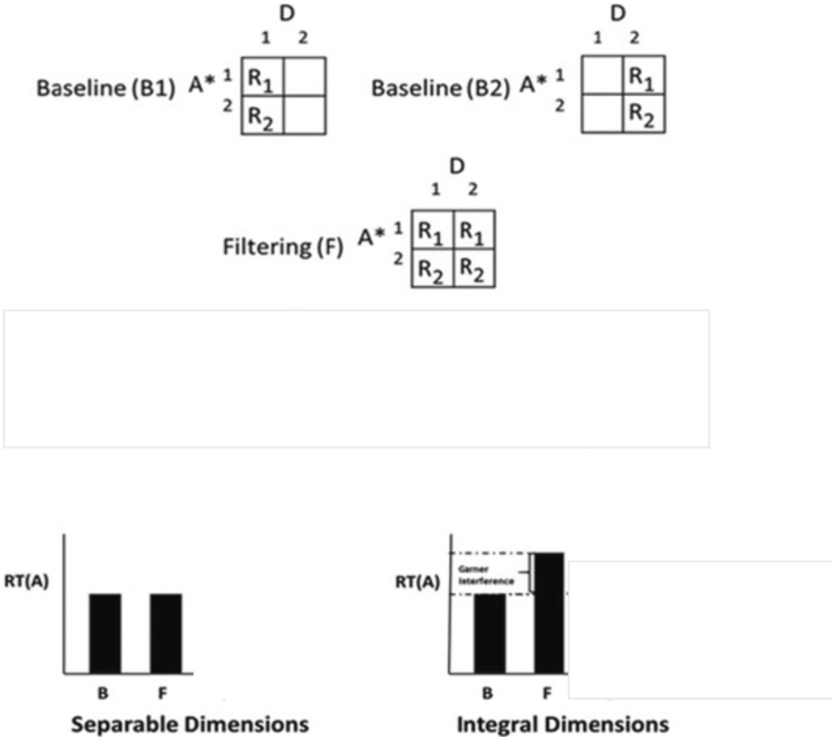

Fig. 2. The Classic Garner paradigm, as illustrated by Algom and Fitousi, 2016, in which participants are required to classify 2D stimuli into two response classes, R_1, R_2. In the Baseline blocks, the task-relevant dimension (A*) varies, while the task-irrelevant dimension (D) is held constant (e.g., Baseline 1 could present wide rectangles and Baseline 2 narrow rectangles, while the height changes in a random fashion from trial-to-trial). In the Filtering blocks, both dimensions vary randomly. If the two dimensions are integral, the failure to focus attention to one while ignoring the other – that is, Garner interference – would be manifested by higher RTs at Filtering as compared to Baseline trials (bottom right graph); such a pattern is not expected to be found when the two dimensions are separable (bottom left graph).

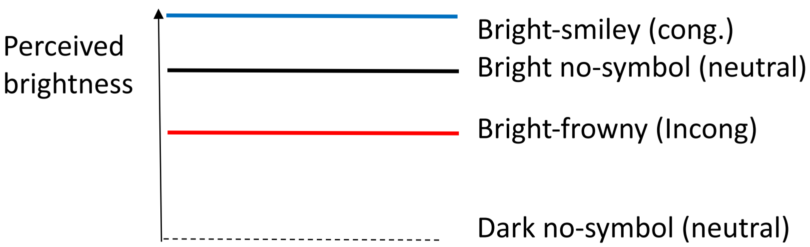

Fig. 3. Illustration of the congruity effect as a result of a modulation of experienced brightness (H2). The y-axis shows the level of experienced brightness, and the two black lines correspond to emotionally neutral gray patches for bright (solid line) and dark (dotted line). The blue line (top) corresponds to a bright patch in a congruent condition (with a smiley symbol), while the red line to a bright patch in the incongruent condition (with a frowny symbol). One can see that the modulation generated by the symbol, increases or decreases the difference between the dark and bright patches. Larger differences (congruent conditions) result is faster RT or higher discrimination levels, compared with small differences (incongruent condition).

congruity effects.

We, therefore, have three possibilities when participants classify the visual brightness of grey patches with induced affective valence as a task-irrelevant dimension:

1. In accordance with H1: The task irrelevant affective valence of the stimulus should neither affect the perceptual experience of brightness nor should it make the stimulus look good/bad (as all alleged affective influences on perceptual experiences are actually the outcome of post-perceptual cognitive appraisals). Hence, we would not expect to find Garner interference. This is because the task is a speeded classification of a low-level perceptual property (brightness), which (in contrasts to tasks involving non-speeded

judgments – see below) is unlikely to be affected by post-perceptual processes. Moreover, Garner interference in particular would support the contention that the affective value of the stimulus is bound with its brightness in a way that requires holistic encoding or processing (Foard & Kemler, 1984; Lockhead, 1972),[10]which is unlikely to take place between perceptual and post-perceptual processes (see General Discussion section).

2. In accordance with H2: The affective value of the stimuli should modulate the perceived sensory properties (the sensory aspect of the experience), making the gray patch look brighter/darker than it is. As illustrated in Fig. 3, this would result in a congruity effect – faster responses for bright-smiley compared to bright-frowny stimuli, and/or a difference in accuracy between congruent and incongruent stimuli.
3. In accordance with H3: The affective value of the stimuli should not alter the perceived sensory properties (thus no congruity effect is predicted),[11] but should introduce a novel component – a valenced aspect – that is strongly bound with (and non-separable from) the sensory dimension. This should result in Garner interference (as for integral dimensions).

To further test the contrast between H1 and H3, we have also set out to examine whether a Garner interference effect takes place at the leading edge of the RT distribution rather than focusing exclusively on mean-RT.[12] Usually, post-perceptual processes are thought to mediate tasks involving non-speeded judgments, notably, via a variety of demand characteristics (Firestone & Scholl, 2016). While our task is a speeded classification task, the aim of this further examination is to reduce even further the likelihood of such post-perceptual contributions. The rationale is that when post-perceptual processes contribute to a task, they add processing time, and thus if Garner interference effect occurs also in the fastest (10 %) of trials, it is very unlikely to involve such processes.

4. Experiment 1: Garner interference with brightness and valence

Participants were required to carry out a speeded classification of brightness for each stimulus, ignoring the presence of superimposed smiley/frowny shapes.

4.1. Method

Participants: A group of 20 university students (mean age = 24.3, SD = 3.2, 14 females), who reported normal color vision, took part in the experiment. This number is the accepted norm in virtually all Garner studies of interference (to illustrate, there were *less* than 20 participants in each of the pair of experiments in the recent Garner study by Niv, Moran, & Algom, 2022).[13] The participants were recruited for credit points. All participants provided written consent for their participation, and the study was approved by the ethical board at Tel-Aviv University.

4.1.1. Materials

The stimuli used in this experiment where (dark/bright) grey circles with superposed smiley/frowny symbols (Fig. 4).

4.2. Procedure and design

Participants were instructed to decide whether the stimulus is bright or dark, as quickly and accurately as possible. The task of the main experiment started with a fixation cross visible for 1 sec. Next, the fixation vanished, followed by a single stimulus presented until answering. Participants pressed “L” for the bright stimuli and “A” for the dark stimuli (no error feedback was presented), and a 500 ms black array was presented after every answer.

Participants completed three blocks – half of the participants did one Filtering block of 80 trials followed by two Baseline blocks of 40 trials each. The second half did the same blocks in reversed order (two Baseline blocks and a Filtering block).

Each type of the four possible stimuli – bright smiley, bright frowny, dark smiley, dark frowny – was chosen in a pseudo-random order, creating two trial types: Congruent trials (smiley bright and frowny dark) and Incongruent trials (smiley dark and frowny bright). In the Baseline blocks, the irrelevant dimension was held constant (the facial expression did not change, only the luminosity). In contrast, the two dimensions varied across trials in a random fashion in the Filtering block.

[10] While binding can take place also between separable dimensions, such binding is fragile and requires attentional effort (Treisman & Schmidt, 1982). By contrast, visual properties from integral dimensions appear to us bound and attention is required to separate them, but this costs time, as reflected in Garner interference. It is thought that this is due to the holistic encoding or perception of the integral dimensions (Fitousi, 2023; Foard & Kemler, 1984; Lockhead, 1972).

[11] Valence and brightness are weak Stroop stimuli, as the former relies on an indirect association, which is unlikely to be automatic. A potential congruity effect would indicate that the association between brightness and our valenced symbols is strong (and fast) enough to affect the classification-RT (for weak Stroop stimuli, see Dishon-Berkovits & Algom, 2000). A null effect, however, would support the expectation that this indirect association is too weak or slow to affect the brightness classification and would also reduce support for the modulation hypothesis, H2.

[12] We formalize this here by comparing the RTs of the 10% fastest responses (10% quantile) in the filtering vs. baseline conditions.

[13] “In response to a reviewer’s request, we estimated using the G-power application (Faul, Erdfelder, Buchner, & Lang, 2009) the statistical power that would be obtained with a sample of N=20 participants, in an experiment with similar effect magnitude as Exp.1 of Niv, Moran, & Algom, 2022. This resulted in a 99.99% power ($\alpha = 0.05$).

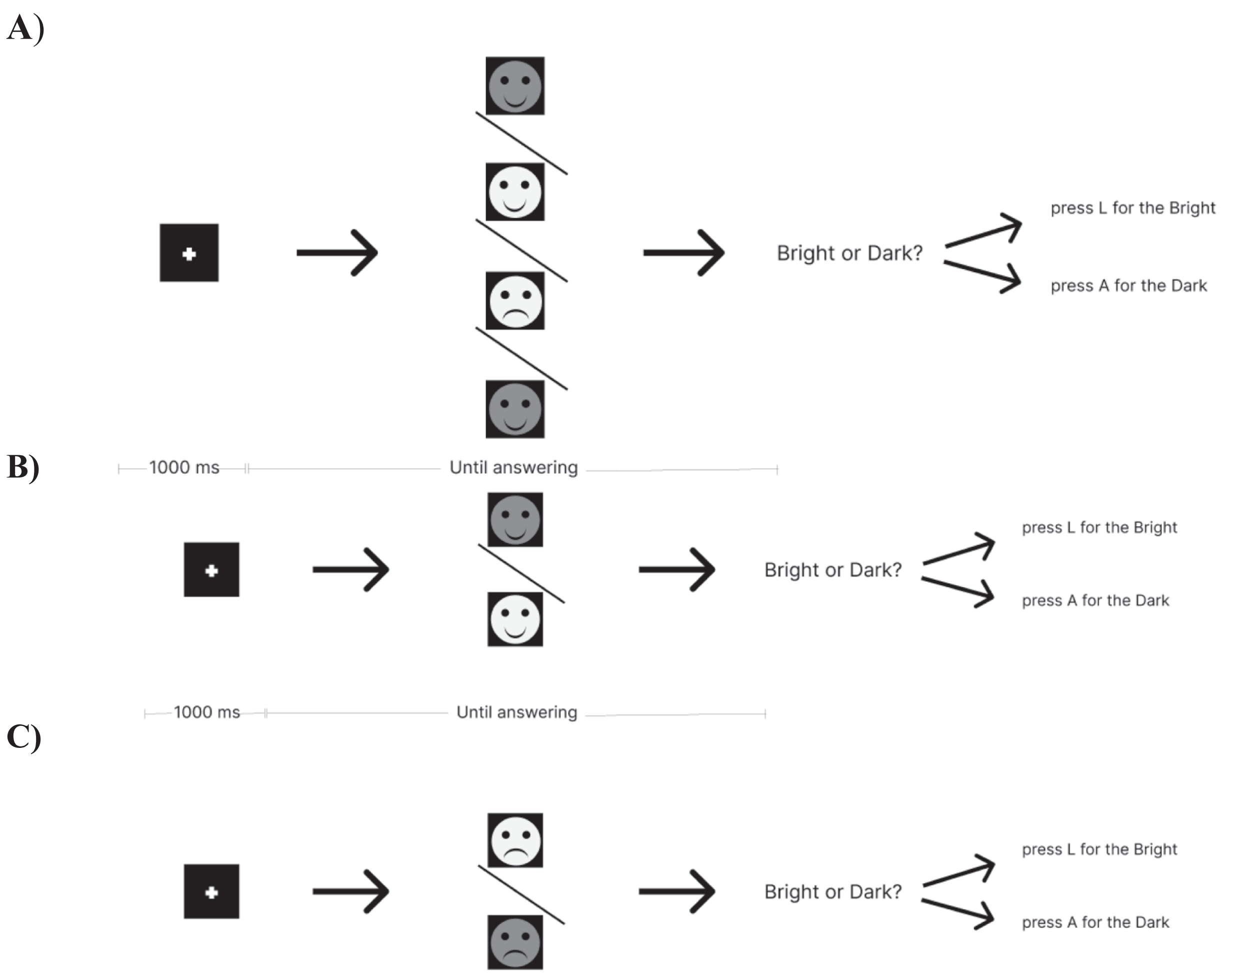

Fig. 4. Illustration of the task structure in the Garner experiment. Three blocks of trials are illustrated, corresponding to filtering and baseline conditions. The bright and dark grey colors differed by an RGB of 90. **A)** Filtering block. Each raw corresponds to a different stimulus type in the filtering condition. Those types of trials are randomly ordered within a filtering block. **B.** Baseline-1 block. (all trials have smiley expressions). The block consists of these two trial types in randomized order. **C.** Baseline-2 block (all trials have frowny expressions). The block consists of these two trial types in randomized order.

When observers failed to respond within 3 s, they were encouraged to respond more quickly, and the trial was marked invalid. A beep sound was played following too slow a response or pressing a wrong key. Aside from that, no feedback was given.

Open Practice: The data and materials for all our experiments can be accessed from: https://osf.io/rxbjy/?view_only=dafbbfd8d45f46fd906a2f957937893b.

4.3. Results

No participant has been removed, as all scored above 90 % accuracy. Mean-RTs for correct responses have been calculated for Filtering (mean-RT = 530 ms; SD = 65 ms) and Baseline (mean = 510 ms; SD = 47 ms). A *t*-test found a significant difference [t(19) = 2.3, p = 0.03; $Cohen'sd = .52$], as shown in Fig. 5A), indicating a significant Garner interference effect of 20 ms.[14] There was no significant difference in accuracy between the two conditions ($t_{(19)} = .59, p > .5, Cohen'sd = .13$).

To detect if the Garner effect we found in this experiment reflects a slow process that is post-perceptual or a fast-perceptual process, we compared the leading edge of the RT distributions. We took the 10 % RT-quantile to measure the leading edge. As shown in Fig. 5B, we find that the difference between filtering and baseline remained highly significant for the 10 % quantile measures [t(19) = 3.62, p = 0.0018; *Cohen's d* = 0.81].

Finally, to assess the potential impact of the Block-type (Baseline vs. Filtering) and Trial-type (Congruent vs. Incongruent) on Mean-

[14] Based on the Cohen-d, this is a medium effect magnitude.

RT, we ran a 2 X 2 within-participant ANOVA for Block-type (filtering/baseline) and Trial-type (congruent/incongruent).[15] We found a significant main effect for Block-type [$F(1,19) = 5.25$, $p = 0.03$; $\eta_p^2 = 0.22$], thus supporting the presence of Garner interference, but not for Trial-type [$F(1,19) = 0.59$, $p = 0.45$; $\eta_p^2 = 0.03$; there was also no congruity effect on accuracy; $p > 0.5$]. The interaction between the two factors was not significant [$F(1,19) = 2.08$,$p = 0.16$;$\eta_p^2 = .1$] (Fig. 6). There was also no significant effect of valence alone: the Mean-RT was not different for trials with smiley (523 ms) and with frowny (516 ms) stimuli $\left(t_{(19)} = 1.01, p = .32, Cohen'sd = .22\right)$.

4.4. Discussion

The presence of Garner interference obtained in Experiment 1 indicates that visual brightness and affective valence are integral dimensions. The participants could not ignore the smiley/frowny symbols when responding to brightness. This result indicates that valence conveyed through these affective symbols is involved in perceptual experience. Note that this interference takes place despite the lack of a valence effect on its own (if at all, responses to negative stimuli were faster than to positive one), unlike in ESE (Algom et al., 2004). Thus, the effect we obtained is specific to Garner interference, which tests how the task varies between blocks (filtering vs baseline), and which specifically probes separable vs. integral perceptual dimensions.

However, one could argue that the Garner effect observed is due to the presence of variation in the task irrelevant shapes independently of their affective association. In other words, the participants took account of the smiley/frowny symbols qua shapes and not as affective stimuli. We believe that this interpretation is unlikely. However, in order to remove any doubt, we have carried out a control experiment, in which we replaced the task-irrelevant smiley/frowny symbols with a new set of symbols (&, $) that lack affective valence.

5. Experiment 2: Task-irrelevant symbols without affective valence

5.1. Methods

Participants: A sample of $N = 20$ participants (same as in Exp. 1; age-range (19–28); 19 female) were recruited from a subject pool at Tel-Aviv University and participated in the experiment for course credit.

5.1.1. Materials and Procedure

As in Experiment 1, except that the smiley/frowny symbols were replaced with ($, &) symbols.

5.2. Results

We compared the mean RT of the correct responses in the Filtering and the Baseline conditions. Unlike in Experiment 1, the RT in Filtering (mean-RT = 525 ms; SD = 68 ms) was not slower than the RT in Baseline (mean-RT = 537 ms; SD = 79 ms). A paired *t*-test showed no significant difference between the two conditions ($t(19) = 0.786$; $p = 0.78$). The main result of Experiment 2 is thus the absence of Garner interference.

The *t*-test is not enough to confirm the null hypothesis (note, however, that the numerical values go opposite to a Garner effect). Therefore, we complemented the test with a Bayesian analysis. We found a $BF_{01} = 7$ in favor of the null hypothesis, suggesting that the observed data are seven times more likely under the null hypothesis than the alternative (presence of Garner interference) hypothesis. This indicates substantial evidence supporting the null hypothesis and hence our interpretation of the results of Experiment 1: the dimensions of visual brightness and affective valence are integral dimensions.

To recap, in Experiment 1, with affective symbols as the task-irrelevant dimension, a sizeable Garner interference was obtained. In Experiment 2, with neutral symbols as the task-irrelevant dimension, we did not observe Garner interference. The difference in outcome indicates the involvement of affective valence in the visual processing and experience of brightness in Experiment 1.

6. General discussion

We set out to test three hypotheses about the effect of induced affective value of a stimulus (positive/negative as manipulated by smiley/frowny symbols) on a basic visual experience: brightness/darkness of grey circles. We considered three hypotheses. The first maintained that induced affective value has no effect on the visual experience, but only on post-perceptual processes (Firestone & Scholl, 2016). The second hypothesis argued that the induced affective value can change the visual experience by making gray circles with smiley symbols appear brighter and gray circles with frowny symbols appear darker (Song et al., 2012; but see Firestone & Scholl 2016). If so, a congruity effect is expected. The third hypothesis was that, independently of a potential sensory modulation as in H2, the affect induced by the smiley/frowny symbols is represented by a valenced (positive/negative) perceptual dimension. Moreover, although valence is a further dimension, orthogonal to brightness, valence and brightness are (according to H3) strongly bound in

[15] In this analysis we averaged (with equal weights) across dark/bright stimuli. Since there were RT differences with stimulus brightness (dark vs bright), we present in the Appendix a 2x2x2 ANOVA, that examines the effects of block and stimulus type for each level of brightness.

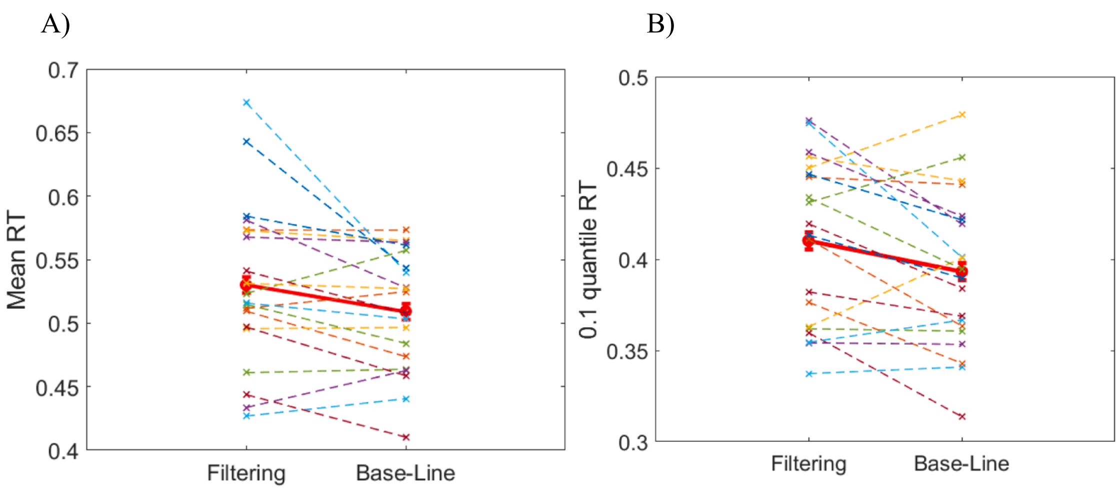

Fig. 5. A) Mean-RT (in seconds) for the correct RT as a function of block-type (filtering/baseline); solid red line with error bars shows the group average, while dotted colored lines show individual participants (error bars correspond to within-subject SE; O'Brien & Cousineeau, 2014). B) Same for 10% RT-quantile.

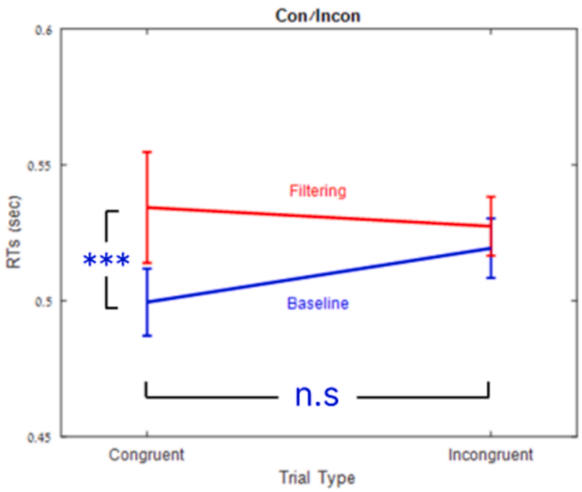

Fig. 6. Mean-RT as function of block (filtering/baseline) and stimulus-type(Cong/Incong).

perception. Thus, valence and sensory aspects jointly determine perceptual performance. This, we suggested, implies that valence and brightness are integral rather than separable dimensions. We thus ran a Garner experiment, in which the task was a binary categorization of brightness, with valence as the irrelevant dimension.

The results of the experiment supported H3. While task accuracy was high (all participants performed at an accuracy rate higher than 90 %), the irrelevant valence produced Garner interference (Fig. 5A). This means that while judging brightness the participants inevitably also registered valence. The results documented that brightness and valence are integral dimensions. Critically, Garner interference was also obtained when we carried out our analysis on the leading edge of the RT-distributions (10 % quantile, which were < 400 ms; Fig. 5B), further supporting a perceptual (rather than post-perceptual) process, consistent with H3. Notably, Garner interference vanished in the control experiment in which we replaced the affective stimuli with non-affective ones. The results, however, did not provide support for H2 (Fig. 3), as there were no congruity effects in the data (see Fig. 6 and Fig. A in the Supplement).[16]

According to H1, the induced affective valence of the smiley/frowny symbols is the outcome of a post-perceptual appraisal. As we have argued, post-perceptual processing is highly unlikely to produce Garner interference. Over and above the basic, low-level task of

[16] While we didn't find support for H2 in our data, it would be incorrect to rule it out from a Null effect. As our focus was on H3 (and H2-3 are not exclusive), we leave the test of H2 for future work.

speeded perceptual classification, the finding that brightness and affective-valence are integral dimensions strengthens the contention that the affective value of the stimulus (together with its brightness) has been perceptually processed. We also contend that, whatever the difference between the formats of perceptual and non-perceptual representations (e.g., iconic format vs. propositional format, respectively), different formats should plausibly make it easier to filter out the perceptual property in speeded classification tasks, and in particular to prevent the dimensions from being bound in the (durable) way that integral dimensions are bound. Lastly, further support for the perceptual interpretation comes from the fact that Garner interference was apparent at the leading edge of the RT distribution, whereas a post-perceptual process would be unlikely to affect the fastest responses.

It should also be noted that while our results did not provide support for H2 in our experimental design, this does not imply that affect cannot result in sensory modulation in other cases. As we have noted, the idea that affect can result in sensory modulations and the idea that perception itself has an affective dimension beyond its sensory dimensions are compatible – i.e., there can be effects of both kinds. Our goal was to argue for the latter idea.

We thus propose that, in at least some cases, perceptual experience is affected by the induced affective value of a stimulus, even in the case of simple sensory discriminations. Furthermore, the impact for which we argue is not one of sensory modulation. Rather, perceptual experiences involve a valenced dimension that is orthogonal to their sensory dimensions. Hence, there is a sense in which things can *look good/bad (positively/negatively)* for the perceiver. Although sensory aspects and valenced aspects are orthogonal, they are not separable dimensions in perceptual experience. Valences and sensory aspects are integral dimensions, such that we cannot selectively attend to one of them while ignoring the other. This is in accordance with Jacobson's (2024) account of the way in which the Valence *in* Perception hypothesis is reflected at the level of appearances, an account that is supported by its theoretical import in resolving various theoretical challenges (such as explaining how valence, as a non-sensory aspect, can be intrinsic to specific sensory experiences). This account appeals to the principle of Intrinsicality as Phenomenal Uniqueness and to the notion of strong, non-compositional, binding: Unlike separable dimensions, specific valences and specific sensory aspects cannot be instantiated independently of each other while maintaining their phenomenal identity – at the level of appearances, they are strongly, non-compositionally, bound together, forming non-factorable looks. The valence changes how a sensory aspect looks, and is thereby 'intrinsic' to a specific look, yet it has no distinct phenomenology independently of the sensory aspect.

While these results show another case of interaction between affect and perception, such as in the Emotional Stroop Effect (ESE), they cannot be explained as a mere slowdown of visual processing due to the simultaneous processing of affect. First, unlike in ESE, we obtained no evidence to show slower responses to frowny compared to smiley stimuli (see section 4). Rather, the effect of emotion depended on how the stimuli were mixed within a block (filtering vs baseline), and it affected both positive and negative valence stimuli, within each block. Second, an emotional slow-down effect would be expected to show in both block conditions (filtering/ baseline), and thus cannot account for such block differences (see Mama et al., 2013, for another demonstration of this principle, based on the contrast between filtering/correlation Garnerian conditions). Rather, the Garner interference is best understood as a failure of attention to separate the task relevant (brightness) from the task-irrelevant stimulus dimensions, which is typically associated with the two dimensions being holistically encoded (Algom et al., 1996).

Our interpretation of the results (as demonstrating that affect is an intrinsic perceptual component) is consistent with previous findings that demonstrate affect to be spontaneously generated (for neutral stimuli) as a result of the facilitation/impediment of perceptual processes (Topolinski et al., 2015; Lindell, Zickfeld & Reber, 2022). However, there are important differences between the two senses of *intrinsicality* implied in these studies. Whereas, in the latter, affect is intrinsic to the perceptual process, in our study, the intrinsicality refers to the perceptual state rather than to the whole process, as formulated by our hypothesis)H3(.

Supporting H3 in the case of valence and brightness is only a small step toward vindicating the Valence *in* Perception hypothesis. The notion that perception has a valenced dimension beyond its sensory dimensions is still in its infancy and faces various challenges. Let us mention just a few of them. First, some crucial open questions concern the encoding of valence as well as the cognitive and neural processes underpinning its binding with sensory aspects. The latter issue is especially important in the context of the present paper, as it raises the question what underlies the fact that the binding of sensory and valenced aspects is a strong, non-compositional binding, or what underlies their forming integral dimensions. This question requires further studies, e.g., on the brain areas that encode valence in and outside the visual cortex and their inter-connectivity, as well as on the receptive fields of the encoding neurons. Though highly speculative, one may consider the possibility that the difference between separable dimensions (e.g., shape and color) and integral ones (hue and brightness for colors, height and width for rectangles, or as we showed here, visual brightness and affective valence) arise from a difference in their brain encoding: distinct brain areas for separable dimensions vs. a single brain area for integral dimensions. Accordingly, the ability to attend to color separately from shape, is related to the distinct brain maps for each of these visual features (Howe et al., 2009; Zeki et al., 1991), and this also accounts for the phenomenological observation of a shared experience when two shapes share a color (Fig. 1A), consistent with mental supervenience on brain states. In the case of integral dimensions, on the other hand, the encoding may take place in a holistic way (e.g., a distributed pattern) in the same brain areas. Obviously, those are theoretical speculations that require empirical validation based on visual neuroscience research.

Second, another pertinent issue concerns how prevalent perceptual valence is (within and across sense-modalities). As mentioned, Lebrecht et al. (2012) and Barrett & Bar (2009) propose that micro-valences characterize visual perception of everyday objects, and it has been suggested that valences are closely associated with basic features such as colors, curved vs. sharp shapes, symmetric vs. asymmetric figures, etc. (e.g., Palmer et al., 2010, Leder et al., 2011). Yet, these suggestions require further support, and even if vindicated, more should be said to support the far-reaching contention that valence is integral to all seeing (Lebrecht et al., 2012; Barrett and Bar 2009; Jacobson 2021).

Third, granted that valenced aspects are not sensory aspects, their very nature is yet unclear. For example, whereas sensory aspects carry information about objective features of objects in our environment, valences do not seem to track objective features, and the affective-values of objects and properties are extremely context-sensitive. Relatedly, these values vary radically both inter-subjectively (i.e., for different subjects) and intra-subjectively (i.e., for the same subject, under different conditions). This has led some to suggest that whereas sensory aspects are clearly representational (or intentional), there are reasons to doubt that the same is true for valenced aspects (see e.g., de Vignemont 2021; Jacobson 2021; for the view that valenced aspects are representational, see e.g., Bain, 2013; 2017). Perhaps, then, valenced aspects should be explained not in terms of evaluative contents – i.e., contents that represent the evaluative properties of the represented items, but rather in terms of evaluative attitudes – i.e., pro/con attitudes directed toward the represented items.

It is worth making explicit the distinction between the question on which we focus in this paper (which the hypotheses address) and the question whether perception is cognitively penetrable (CP). The latter question concerns the possibility that visual experiences are not encapsulated from higher-level cognition; rather, cognitive factors, including affective factors, "exert direct, top-down influences on what we see" (Firestone and Scholl, 2016). We inquire whether, due to the affective value of stimuli, perception can have an integral valenced dimension and/or whether the affective value of stimuli can modulate the sensory properties of perception. Suppose the answer is positive. In this case, there is a further question – the CP question, which is concerned with *how* – i.e., by what means or mechanisms – the affective value of a stimulus influences perception. Focusing for illustration on the Valence *in* Perception hypothesis, if perception is shown to have an integral valenced dimension, the question arises *how* perception *comes to have* that valenced dimension. One possibility is that the affective value of the stimulus first forms in us a (higher-level) valenced cognitive state (e.g., a mediating thought, expectation, or emotion with an evaluative content, such as 'the stimulus is bad to some degree') that then influences the perceptual experience, by making it the case that the stimulus looks to us bad to some degree. This would be an instance of CP. Another possibility, which we take to be quite plausible in certain central cases, is that the relevant process involves no (direct) top-down influences; rather, the (early processed) affective value of the stimulus is itself directly perceived. That is, it is entirely coherent to adopt H2 and/or H3 while rejecting the notion of CP. As such, we will remain agnostic regarding the question of CP.

The difficulties surrounding the Valence *in* Perception hypothesis notwithstanding, this hypothesis potentially has exciting implications for the understanding of the function and the significance of conscious experience. The idea that conscious experiences are inherently affective suggests that "perception involves, along with its more traditional role, an additional evaluative and motivating function" (Fulkerson, 2020). Whereas the sensory aspects of perceptual experiences are geared toward providing accurate information about objective features of the environment, the role of valenced aspects plausibly has to do with evaluating those features in relation to the subject's specific needs and present states. Relatedly, the affective value of stimuli plausibly plays a role in motivating or inhibiting action toward them and guiding rudimentary kinds of selection among them. Such functions are pertinent to survival, and there are plausibly processing advantages to their being integral to perceptual processes. Hence, this dual conception of the function of perception which the Valence *in* Perception hypothesis affords renders this hypothesis attractive from an evolutionary perspective.

Last but not least, phenomenal consciousness matters – there is a strong intuition that the fact that an organism has conscious experiences implies that it has some level of well-being (whether positive or negative) and renders its life (if only to some degree) significant. Yet, if perceptual experiences are 'cold' and affectless, why should consciousness, as such, be significant? Why should we care whether an organism has conscious experiences? If the Valence *in* Perception hypothesis is correct, this conundrum rests on a mistake: consciousness might be fundamentally evaluative and valenced, and as such it matters to the conscious subject and merits some level of concern from others.

CRediT authorship contribution statement

Hilla Jacobson: Writing – review & editing, Writing – original draft, Conceptualization. **Zohar Ongil:** Writing – review & editing, Software, Investigation, Data curation. **Daniel Algom:** Writing – review & editing, Investigation, Formal analysis. **Marius Usher:** Writing – review & editing, Writing – original draft, Supervision, Project administration, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We wish to thank Dikla Sacks for testing participants in Experiment 2 and for data analysis, Gal Atun for data analysis, Moshe Glickman for statistical assistance and Niccolo Negro for a critical reading of the manuscript. HJ was supported by an Israel Science Foundation grant (ISF-1083-22). MU was supported by an Israel Science Foundation grant (ISF-1602-22). DA was supported by an Israel Science Foundation grant (ISF-543-19).

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.concog.2024.103783.

Data availability

Data will be made available on request.

References

Algom, D., Dekel, A., & Pansky, A. (1996). The perception of number from the separability of the stimulus: The Stroop effect revisited. *Memory & Cognition, 24*, 557–572.
Algom, D., Chajut, E., & Lev, S. (2004). A rational look at the emotional Stroop phenomenon: A generic slowdown, not a Stroop effect. *Journal of Experimental Psychology: General, 133*, 323–338.
Algom, D., & Fitousi, D. (2016). Half of a century of research on Garner interference and the separability-integrality divide. *Psychological Bulletin, 142*, 1352–1383.
Bain, D. (2013). What Makes Pains Unpleasant? *Philosophical Studies, 166*(1), 69–89.
Bain, D. (2017). Why Take Painkillers? *Noûs, 53*(2), 462–490.
Banerjee, P., Chatterjee, P., & Sinha, J. (2012). Is it light or dark? Recalling moral behavior changes perception of brightness. *Psychological Science, 23*, 407–409.
Barbot, A., & Carrasco, M. (2018). Emotion and Anxiety Potentiate the Way Attention Alters Visual Appearance. *Scientific Reports, 8*(1), 1–10.
Bargh, J. A. (1992). The ecology of automaticity: Towards specifying the conditions necessary to produce automatic processing effects. *American Journal of Psychology, 105*, 181–199.
Bargh, J. A., Chaiken, S., Govender, R., & Pratto, F. (1992). The generality of the automatic attitude activation effect. *Journal of Personality and Social Psychology, 62*, 893–912.
Barrett, L. F., & Bar, M. (2009). See it with feeling: Affective predictions during object perception. *Philosophical Transactions of the Royal Society B: Biological Sciences, 364*(1521), 1325–1334.
Birch, L. (1999). Development of Food Preferences. *Annual Review of Nutrition, 19*, 41–62.
Block, N. (2023). *The Border between Seeing and Thinking*. New York, US: OUP USA.
Carrasco, M., & Barbot, A. (2019). Spatial Attention Alters Visual Appearance. *Current Opinion in Psychology, 29*, 56–64.
Chen, M., & Bargh, J. A. (1999). Consequences of automatic evaluation: Immediate behavioral predispositions to approach or avoid the stimulus. *Personality and Social Psychology Bulletin, 25*, 215–224.
Cohen J, Dunbar & McClelland J (1990). Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97* (3), 332–361.
De Houwer, J., Thomas, S., & Baeyens, F. (2001). Associative Learning of Likes and Dislikes: A Review of 25 Years of Research on Human Evaluative Conditioning. *Psychological Bulletin, 127*(6), 853–869.
de Vignemont, F. (2021). Fifty shades of affective colouring of perception. *Australasian Journal of Philosophy*, 1–15.
Dishon-Berkovits, M., & Algom, D. (2000). The Stroop effect: It is not the robust phenomenon that you have thought it to be. *Memory & Cognition, 28*, 1437–1449.
Domínguez-Borràs, J., Rieger, S. W., Corradi-Dell'Acqua, C., Neveu, R., & Vuilleumier, P. (2017). Fear spreading across senses: Visual emotional events alter cortical responses to touch, audition, and vision. *Cerebral Cortex, 27*(1), 68–82.
Duncan, S., & Barrett, L. F. (2007). Affect is a form of cognition: A neurobiological analysis. *Cognition and Emotion, 21*(6), 1184–1211.
Eidels, A. (2012). Independent race of colour and word can predict the Stroop effect. *Australian J. of Psychology, 64*, 189–198.
Faul F, Erdfelder E, Buchner A, & Lang AG (2009). *Behavior Research Methods 2009, 41 (4), 1149-1160.*.
Ferneyhough, E., Kim, M. K., Phelps, E. A., & Carrasco, M. (2013). Anxiety modulates the effects of emotion and attention on early vision. *Cognition & Emotion, 27*, 166–176.
Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behavioral and Brain Sciences, 39*, e229.
Fitousi, D. (2023). Dimensional interaction is driven by feature integration: A binding theory of Garner interference. *Journal if Experimental Psychology: Learning, Memory, and Cognition, 49*, 1881–1908.
Foard, C. F., & Kemler, D. G. (1984). Holistic and analytic modes of processing: The multiple determinants of perceptual analysis. *Journal of Experimental Psychology: General, 113*(1), 94–111.
Freud, E., & Ganel, T. (2015). Visual control of action directed toward two-dimensional objects relies of holistic processing of object shape. *Psychological Bulletin & Review, 22*, 1377–1382.
Fulkerson, M. (2020). Emotional Perception. *Australasian Journal of Philosophy, 98*(1), 16–30.
Ganel, T., & Goodale, M. A. (2014). Variability-based Garner interference for perceptual estimations but not for grasping. *Experimental Brain Research, 232*, 1751–1758.
Garner, W. R. (1974). *The processing of information and structure*. Potomac, MD: Erlbaum.
Handy TC, Smilek D, GeCiger L, Liu, & Schooler JW (2010) ERP evidence for rapid hedonic evaluation of logos. *Journal of Cognitive Neuroscience, 22*, 124-138.
Howe PDL, Evans KK, Pedersini R, Horowitz TS, Wolfe JM, & Cohen MA (2009). Attention: Selective Attention and Consciousness, Ed. William P. Banks, Encyclopedia of Consciousness, 61-75.
Jacobson, H. (2021). The Role of Valence in Perception: An ARTistic Treatment. *Philosophical Review, 130*(4), 481–531.
Jacobson, H. (2024). On the Very Idea of Valenced Perception. *Journal of Philosophy*. in press.
Kuhbandner, C., Spitzer, B., & Pekrun, R. (2011). Read-out of emotional information from iconic memory: The longevity of threatening stimuli. *Psychological Science, 22*, 695–700.
Lebrecht, S. (2012). *Micro-valences: Affective Valence in 'Neutral' Every- day Objects*. PhD diss.: Brown University.
Lebrecht, S., & Tarr, M. (2010). Defining an Object's Micro-valence through Implicit Measures. *Journal of Vision, 10*(7), 966.
Lebrecht, S., Bar, M., Barrett, L. F., & Tarr, M. J. (2012). Micro-valences: Perceiving affective valence in everyday objects. *Frontiers in psychology, 3*, 107.
Leder, H., Tinio, P. L., & Bar, M. (2011). Emotional Valence Modulates the Preference for Curved Objects. *Perception, 40*(6), 649–655.
LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience, 23*, 155–184.
Lindell, T. A. E., Zickfeld, J. H., & Reber, R. (2022). The Role of Affect in Late Perceptual Processes: Evidence from Bi-Stable Illusions, Object Identification, and Mental Rotation. *J. of Exp. Psychol. HPP, 48*(12), 1347–1361.
Lockhead, G. R. (1972). Processing dimensional stimuli: A note. *Psychological Review, 79*(5), 410–419.
MacLeod, C. M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin, 109*, 163–203.
Mama, Y., Ben Haim, M. S., & Algom, D. (2013). When emotion does and does not impair performance: A Garner theory of the emotional Stroop effect. *Cognition & Emotion, 27*(4), 589–602.
McKenna, F. P. (1986). Effects of unattended emotional stimuli on color-naming performance. *Current Psychology, 5*, 3–9.
Meier, B. P., Robinson, M. D., Crawford, L. E. & Ahlvers, W. J. (2007) When "light" and "dark" thoughts become light and dark responses: Affect biases brightness judgments. *Emotion* 7(2):366–76. Available at: http://doi.org/10.1037/1528- 3542.7.2.366.
Melara, R. D., & Algom, D. (2003). Driven by information; A tectonic theory of Stroop effects. *Psychological Review, 110*, 422–471.
O'Brien, F., & Cousineau, D. (2014). Representing error bars in within-subject designs in typical software packages. *The quantitative methods for psychology, 10*(1), 56–67.
Palmer, S. E., & Schloss, K. B. (2010). An Ecological Valence Theory of Human Color Preference. *Proceedings of the National Academy of Sciences, 107*(19), 8877–8882.
Pessoa, L. (2008). On the relationship between emotion and cognition. *Nature reviews neuroscience, 9*(2), 148–158.
Phelps, E. A., Ling, S., & Carrasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological Science, 17*, 292–299.

Pourtois, G., Schettino, A., & Vuilleumier, P. (2013). Brain mechanisms for emotional influences on perception and attention: What is magic and what is not. *Biological Psychology, 92*, 492–512.
Reber, R., Winkielman, P., & Schwartz, N. (1998). Effect of perceptual fluency on affective judgments. *Psychological Science, 9*, 45–48.
Niv, L., Moran, R., & Algom, D. (2022). The nature of Garner interference: The role of uncertainty, information, and variation in the breakdown in selective attention. Cognition. Jan;218:104950. doi: 10.1016/j.cognition.2021.104950.
Reber, R., Schwartz, N. & Winkielman, P. (2004) Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience? *Personality and Social Psychology Review* l8:364–382.
Sabri, M., Melara, R. D., & Algom, D. (2001). A confluence of contexts: Asymmetric versus global failures of selective attention to Stroop dimensions. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 515–537.
Song, H., Vonasch, A. J., Meier, B. P., & Bargh, J. A. (2012). Brighten up: Smiles facilitate perceptual judgment of facial lightness. *Journal of Experimental Social Psychology, 48*, 450–452.
Stefanucci, J. K., & Storbeck, J. J. (2009). Don't look down: Emotional arousal elevates height perception. *Journal of Experimental Psychology: General, 138*, 131–145.
Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology, 18*, 643–662.
Topolinski, S., Erle, T. M., & Reber, R. (2015). Necker's smile: Immediate affective consequences of early perceptual processes. *Cognition, 140*, 1–13.
Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14*(1), 107–141.
Vuilleumier, P. (2015). Affective and Motivational Control of Vision. *Current Opinion in Neurology, 28*(1), 29–35.
Williams, J. M. G., Mathews, A., & MacLeod, C. (1996). The emotional Stroop task and psychopathology. *Psychological Bulletin, 120*(1), 3–24.
Zeki, S., Watson, J. D., Lueck, C. J., Friston, K. J., Kennard, C., & Frackowiak, R. S. (1991). A direct demonstration of functional specialization in human visual cortex. *J Neurosci., 11*(3), 641–649.